

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE Nov 9, 94	3. REPORT TYPE AND DATES COVERED Final 7/1/94-11/1/94		
4. TITLE AND SUBTITLE Summer Institute on Probability in AI Final Report		5. FUNDING NUMBERS F49620-94-1-0348 230416S 01102F		
6. AUTHOR(S) Bruce D'Ambrosio		8. PERFORMING ORGANIZATION REPORT NUMBER AFOSR-TR-95-0277		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Oregon State University Computer Science 303 Dearborn Hall Corvallis, OR 97331-3202		10. SPONSORING/MONITORING AGENCY REPORT NUMBER F49620-94-1-0348		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NM		11. SUPPLEMENTARY NOTES DTIC SELECTED JUN 19 1995 F		
12a. DISTRIBUTION/AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) The first Summer Institute on Probability in AI was held July 22 - 29, 1994 on the campus of Oregon State University. This report reviews the motivation for and organization of that workshop, and also briefly reviews the current status of the field, as revealed at the workshop. DTIC QUALITY INSPECTED 5				
14. SUBJECT TERMS		15. NUMBER OF PAGES 20		
		16. PRICE CODE		
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT SAR	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

19950615 044

Summer Institute on Probability in AI Final Report

Bruce D'Ambrosio
Department of Computer Science
Oregon State University
Corvallis, OR
dambrosi@research.cs.orst.edu

November 7, 1994

Accession For	
NTIS CRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution /	
Availability Code	
Dist	Avail and/or Special
A-1	

Abstract

The first Summer Institute on Probability in AI was held July 22 - 29, 1994, on the campus of Oregon State University in Corvallis, Oregon. This report reviews the motivation for and organization of that workshop. It also briefly reviews the current status of the field, as revealed at the workshop.

1 Workshop objectives

The institute was designed as an intensive review of modern Bayesian modeling and inference techniques, focusing on belief networks and influence diagrams. The idea was to collect in one place many of the leading researchers in the field together with advanced Phd students, recent Phds, and government and industrial researchers who might wish to apply these methods. It was hoped the resulting interchange would help spread the word about rapid developments in the field, recruit new researchers to the area, and help define needed research directions.

Recent developments in Bayesian modeling, both representational and inferential, (i.e. Bayes nets and associated algorithms) represent a significant breakthrough in representation of and inference with uncertain information. However, there is a significant startup cost to incorporating these developments into ongoing research programs. This has been a substantial inhibitor to the spread of these techniques. Thus, one motivation for the workshop was to widely distribute "seeds of expertise" who would foster the dissemination of knowledge about the technology. At the same time, research in some areas of Bayesian Nets, eg inference, was beginning to reach a plateau. It was clear that researchers need an infusion of fresh problems to help define research directions and focus research. Thus, a second objective was to bring leading researchers in close contact with domain experts from a variety of fields who could provide new research challenges.

2 Workshop organization

We felt these objectives could best be achieved through an intensive summer workshop or institute which gathered in one place the leaders in the field with the best current graduate students and interested application-oriented researchers. Such summer schools have been quite successful in other areas (for example, the connectionist summer schools held in the late 80's). We developed an outline for a Bayesian modeling summer school, recruited faculty, and obtained commitments from community members to serve on various organizing committees (Special thanks to Andrew Mayer of Heuristicrats Research, who managed the participant application process). The syllabus was developed interactively through email discussions involving most of the U.S. Bayes net community. We anticipate that the summer school ran for seven days, and was a mix of lectures and "hands-on" labs. It began with foundational discussions of modeling and knowledge acquisition, continued with a review of reasoning methods, and concluded with an overview of advanced applications in diagnosis, robotics, vision. The final schedule, which includes abstracts for all presentations and labs, is included as appendix I. The materials distributed at the workshop are too extensive to include with this report, but one complete set has been submitted to AFOSR. One complete set of the handouts and slides has been sent to

2.1 Participation

Interest in the workshop was quite high. Despite very limited advertising, we received over 160 applications to attend. We chose participants based on a desire for diversity,

geographical, technical, and organizational. Participants included handwriting and natural language experts from Calera and Sun Microsystems, botanists and geographers from U Missouri and UC Santa Barabra, and atr and diagnosis experts from Wright Labs and Nasa Marshall. A complete list of attendees is given in appendix II. DOD related participants constituted a significant percentage of those attending. Table 1 shows a list of the DOD and NASA affiliated participants.

Name	Home Institution	DOD contract
Greg Arnold	Wright-Patterson AFB	
Marco Barbarisi	US Navy Coastal Sys Div.	
Michael Frank	MIT	ARPA/Rome Labs Planning Initiative F30602-91-C-0018
Jordan Hayes	Heuristicrats Research	Warbreaker
Michael Jordan	MIT	ONR N00014-90-J-1942
Jim Leonard	Wright-Patterson AFB	
Kurt McCall	NASA Marshall Elec. Pwr	
Rulane Merz	Hughes Information Tech	ARPA F33657-94-C-4073.
Mark Minardi	Wright-Patterson AFB	
Stan Musick	Wright-Patterson Labs	
Gary Ogasawara	Lockheed AI Center	USN Coastal Systems N61331-91-C-0028
Barney Pell	NASA Ames AI Branch	
Mehdi Shirazi	Wright-Patterson AFB	
Suzanne Stanford	George Mason Univ. CS	Warbreaker
Irena Vainshtein	Litton Data Systems	ARPA/MSTO: SEAPACS
Tony Zawilski	Mitre	??

Figure 1: DOD and NASA affiliated Participants

3 Report from the Summer Institute: the State of the Art

This section will review the summer institute from two perspectives. First, it will review the state of the art in Bayesian networks, as revealed at the institute. Second, it will highlight important developments, new opportunities, and outstanding research issues revealed at the workshop. The review will follow the topic organization and ordering used at the institute itself.

3.1 Foundations

Clarity was perhaps the most difficult foundational issue for institute participants. Clarity is a requirement placed on all variables in a probabilistic model, Bayesian or otherwise. Informally, the clarity requirement is satisfied if one can imagine that an oracle could

determine categorically the single true value of the variable. The meaning of and the need for clarity were both difficult for many participants to understand.

3.2 Knowledge Acquisition

Two varieties of knowledge acquisition were discussed at the institute, standard discrete Bayes net construction and mixed discrete/continuous modeling. Participants were provided the opportunity to explore model construction in two labs, and open labs were available throughout the remainder of the institute to permit participants to explore the available tools further.

Assessment: Knowledge acquisition methods for Bayesian networks are reasonably mature. However, it remains an arduous task, similar to but more well understood than knowledge acquisition for other forms of knowledge-based systems. One area of potential further progress is the identification of new canonical interaction models (eg, noisy or).

3.3 Inference

The inference session at the institute was divided into two parts, basics and research topics. As with knowledge acquisition, the basics seem well understood. A perspective growing in popularity is the view of efficient inference as fundamentally a rewrite problem to obtain efficiently evaluable expressions for distributions of interest. While it is unlikely that this will result in algorithms with dramatically improved performance, it may lead to greater insight into algorithm performance and to simpler, easier to implement, algorithms.

Research topics included extended representations and approximate evaluation. Research on extended representations, and efficient algorithms for them, is in a period of rapid progress. This progress is needed: existing exact algorithms are computationally intractable for large classes of practical problems, and existing approximation algorithms (primarily monte-carlo based) are of limited performance.

Assessment: Exact inference for standard belief nets is well understood. Further refinement of simulation (monte-carlo) based approximation schemes is unlikely to produce dramatic performance improvements. Yet many problems of practical interest are beyond the reach of these methods. Extended representations, capable of capturing intra-distribution structure, together with algorithm extensions capable of exploiting this structure, offer one possibility for overcoming this computational barrier, and deserve serious study. Novel approximation methods, such as search-based methods and successive representation refinement methods offer an alternative promising path.

3.4 Applications

The Applications portion of the institute began with the most tractable application, diagnosis, and proceeded to more difficult applications, including language understanding, vision, and planning. However, it needed have gone further than diagnosis to find challenges: The diagnosis discussions made it clear that careful engineering of the representation is necessary to keep most diagnostic applications tractable. Diagnosis of static

systems (ie, system state does not change during diagnosis) under a single fault assumption is well understood, practical, and widely applicable. Problems of this class should see increasingly wide-spread use of Bayesian methods in future. However, multiple-fault, and especially dynamic systems (ie, systems which can change state during diagnosis) present significant challenges for future research.

The more advanced applications present even greater challenges to general purpose methods. The reviews of natural language, vision, and planning made two things clear: (1) Bayesian methods in general, and Bayesian networks in particular, have major contributions to make to these areas; (2) current methods, and even current problem definitions, within the Bayesian network algorithms community are inadequate for these tasks. Effective support will require breaking down the barriers between inference algorithms and the embedding application. More all-encompassing task definitions, including model extension, model reformulation, and resource allocation within the scope of task definition, are needed to support these applications.

Assessment: Some applications (single fault diagnosis, some model-based vision applications) are practical now. Other, more advanced applications will require concerted effort to generate advanced, integrated methods.

3.5 Time and causality

Earlier in this report we stated that the foundations of belief nets were well understood. The sessions on time and causality made it clear that this is true only for static or discrete-time, acausal models. It should be no surprise to discover that all the problems of representing time in logical approaches to representation recur in probabilistic approaches. It may be a surprise, however, to realize that the semantics of causality in Bayesian Networks, even those assessed in the "causal" direction, are unclear. However, all of the issues around the distinction between correlation and causality arise in Bayes net representations, and these become particularly relevant when one attempts to use them as a basis for planning. The distinction between a chance node and a decision node in an influence diagram is in part an assertion of causality: If I change the value of the decision node, the value of the consequent nodes will change. But in an influence diagram this is confounded with another assertion: I must decide what value to assign to this node.

Assessment: Many more mundane applications are practical now, but advancement of the longer term goals of AI and intelligent robotics will demand progress in understanding of the issues of time and causality.

3.6 Qualitative models

The institute revealed that Qualitative models of probability is a very active research area at present. System Z (Pearl, et al) is being explored by a number of research groups. Alternate formulations, such as abstraction of variable state spaces, are also attracting interest.

Assessment: As mentioned earlier, existing exact inference methods have serious limitations in many applications. It is the hope of many that qualitative methods will

overcome these limitations, but further research is needed to see if that hope will be fulfilled.

3.7 Learning

Bayesian methods, at least at the meta-level, are fast becoming predominate in machine learning. Direct implementations of Bayesian methods also seem quite effective at the performance level. Bayesian methods are the methods of choice when available data is limited, and methods for learning Bayes nets are beginning to move out of the laboratory.

Assessment: Current methods seem quite effective when very strong constraints on the overall structure of the model are available (eg, all parameters known, and a total ordering on the parameters is given). Newer methods are finding ways to relax these requirement slightly. This seems an area of rapid progress and great potential.

3.8 Model Confidence and Resource Bounding

Earlier in this report we stated that exact inference was well understood. However, the task definition is too small, as was seen in the discussion of advanced applications. Model Confidence and resource-bounded reasoning issues both arise when we seek to expand the boundaries of the inference task. Model confidence methods attempt to answer the question: given evidence, can we draw any conclusions about the adequacy of the model we are using? Resource-bounding methods attempt to balance quality of solution with computational resources expended.

Assessment: Solid theoretical foundations have been laid in both areas, but little practical application work has been done. It remains to be seen how much work in this area will extend the boundaries for direct application of formal probabilistic methods.

4 Summary

The First Summer Institute on Probability in AI was held between July 22 and July 27, 1994, in Corvallis, Or. Fifteen speakers and 55 non-speakers attended, including fourteen with DOD affiliations. The sessions at the Institute confirmed that, while a great amount remains to be done, a significant technology is now available and ready for practical application in a variety of domains. At the same time, specific research problems were identified in the areas of inference, learning, and task definition for advanced applications.

5 Appendix I - Schedule

Friday, 7/22

AM: 7:00 - 8:00 Breakfast and Registration, Gold Room, Arnold

AM: 8:30 - 8:35 Opening Remarks, Bruce D'Ambrosio - ECE 102

8:35 - 8:45 Welcome, Abraham Waksman, AFOSR

8:45 - 12:00 Foundations

Ross Shachter, Stanford

This session will address the assumptions and philosophy underlying Bayesian probabilistic and decision analysis models. We define conditional independence and explore some of its properties. There will be a large number of examples of belief networks and influence diagrams without numbers, to build an understanding of these models on a structural level. This session will be as interactive as possible, with the understanding that we have a number of concepts to review.

AM: 10:15 - 10:30 Break, ECE Lobby

12:00 - 1:30 Lunch, Gold Room, Arnold

PM: 1:30 - 5:00 Knowledge Acquisition/Model Construction Bexel 417
Jack Breese, Microsoft Research
David Heckerman, Microsoft Research
Max Henrion, Lumina Decision Systems

In this section we will cover a number of issues related to model formulation, probability assessment, and preference assessment in the context of belief networks and influence diagrams. Presentations will be interleaved with two labs, one emphasizing predictive modeling, the other construction of belief networks for diagnostic applications.

A summary of the topics to be covered:

Selecting the best approach

- Normative, descriptive, heuristic, and prescriptive methods
- Rule-based vs decision-theoretic schemes
- Prediction versus diagnosis

Structuring knowledge with IDs

- Problem formulation
- Defining the application
- Distinguishing decisions, chances, and objectives
- Evaluating and identifying tradeoffs.

Quantifying probability distributions

- Discrete probability assessment
- Continuous probability assessment
- Calibration, biases and heuristics
- Prototypical influences and the noisy-OR

Eliciting preferences

- Utility assessment
- Risk attitudes
- Multiattribute utility functions

Iterative Modeling

- Evaluating influence diagrams
- The value of information
- The decision analysis cycle and sensitivity analysis

PM: 2:45 - 3:00 Break, Bexell 417

PM: 3:00 - 5:00 KA/MC Lab, CS West

Saturday, 7/23

AM: 7:00 - 8:00 Breakfast, Gold Room, Arnold

AM: 8:30 - 12:00 Continuation of KA/Model Construction, Bexell 417

AM: 10:00 - 10:15 Break, Bexell 417

AM: 10:30 - 12:00 KA/MC lab, Bexell 228

12:00 - 1:30 Lunch, Gold Room, Arnold

PM: 1:30 - 5:00 Algorithms 102 ECE
Robert Fung, Lumina Decision Systems
Bruce D'Ambrosio, Oregon State University

In this talk we presume theoretical foundations have been laid by Ross Shachter, and discuss some of the more "practical" issues around inference. We begin with a review of the basic inference tasks (eg, prediction, posterior computation, and optimal policy determination) and algorithm classes. We then focus on two or three exact algorithms and examine methods for importance sampling. The first half of the talk closes with a review of available tools. The second half of the talk then discusses current research in algorithms, including algorithms for noisy or and other advanced representations, and approximate algorithms including search and simulation-based approaches.

PM: 3:00 - 3:15 Break

Dinner: 7:00 - 9:00, Gold Room, Arnold
Perspective: Normative Systems, Ward Edwards, USC

Sunday, 7/24 - Free day.

Brunch, 10am - 11am, Gold Room, Arnold.

Note - No Breakfast or Dinner, you are on your own!

Bexell lab open 1pm - 5pm

CS lab open 8am - midnight

Monday, 7/25

AM: 7:00 - 8:00 Breakfast: Gold Room, Arnold

AM: 8:30 - 12:00 Diagnosis 103 ECE

Mark Peot, Knowledge Industries

Greg Provan, Institute for Decision Systems Research

This talk will cover both the theoretical and the practical side of diagnosis using influence diagrams. Approximately half of the talk will be a general overview of diagnostic reasoning. The other half consists of the discussion of several case studies drawn from the domains of medical and machine diagnosis.

Topics to be covered:

GENERAL OVERVIEW

- * Description of the Diagnosis Task
- * Computational Properties of Diagnosis
- * Mechanisms for Making Diagnosis Practical
 - Value of Information Techniques
 - Simplified Diagnosis Models
- * Knowledge Representation for Diagnosis
 - Types of Models (Single Fault, Multiple Fault)
 - Deep vs Shallow Models
 - Dependency (Full Dependency vs Independence Assumptions)
 - Distribution Types (Kappa, Point Probabilities, Density Functions)
- * Diagnostic Focussing
 - Knowledge-Based Model Construction
 - Submodel Selection

Case Studies

- * Medical Diagnosis: QMR-DT/CPCS, Acute Abdominal Pain
- * Machine Diagnosis: Garrett 700 APU, NASA Vista

AM: 10:15 - 10:30 Break ECE lobby

12:00 - 1:30 Lunch: Gold Room, Arnold

PM: 1:30 - 3:00 Natural Language 103 ECE
 Robert Goldman, Honeywell SRC

We are concerned with the problem of natural language understanding, viewed as probabilistic inference. We see natural language understanding as an abduction, or diagnosis, problem. One is given an utterance, or a text, and would like to "diagnose" the intent of the person who produced that utterance, where that intent is usually (but not always) to communicate some piece of information. We assume that probabilistic inference is the Right Thing for solving such abduction problems.

In-depth understanding of natural language utterances [as opposed to e.g., parsing, morphological analysis, stemming...] presents some interesting difficulties for graphical modeling. There does not seem to be a useful graphical model for "the utterance," or "the sentence," rather there is a class of problems, each utterance being one member of this class. Issues of abstraction are crucial: a vague, high-probability interpretation is considerably less interesting than a somewhat less likely interpretation which contains specific information. There is beginning to be interesting statistical information about natural language utterances, but only at surface levels. Interpretations still require engineering. How can we exploit both subjective and empirical estimates?

In this lecture we will provide a conceptual outline of the problem of natural language understanding, present an approach to this problem based on knowledge-based model construction and present some interesting open research issues.

PM: 3:00 - 3:15 Break

PM: 3:15 - 5:00 Vision 103 ECE
Tod Levitt, IET, Inc. and Stanford University

Vision abstract:

We demonstrate that a belief system, and Bayesian inference and decision theory in particular, are critical to successful automated machine understanding of imagery in general. The following topics are explored:

- * definitions and mechanisms for computer vision
- * complexity and sources of uncertainty in image understanding
- * generic versus specific models and their inferential interpretations
- * model-based vision and the meaning of Bayesian inference in it
- * evidence and the impact of its computation
- * hierarchies of reasoning, inference and computation
- * parallel processing, parallel hypotheses and inference
- * quasi-invariants and the semantics of observational induction

Machine or computer vision is usually cited as an application of uncertain reasoning, but in the "perceive, think, act" view of robotic intelligence, computer vision is as fundamental a technology in AI as automated planning or diagnosis, which are usually viewed as incarnations of machine inference, rather than applications. This somewhat controversial view suggests that a great emphasis should be placed on machine perception by the AI and Uncertainty in AI communities as the inevitable path to widespread real world applications of uncertainty reasoning in machine systems.

PM: 7:00 - 8:00 Dinner: Gold room, Arnold,

Bexell lab open 8am - 5pm and 7pm - 10 pm
CS lab open 8am - midnight

Tuesday 7/26

AM: 7:00 - 8:00 Breakfast: Gold Room, Arnold

AM: 8:30 - 9:45 Temporal Reasoning and Representation - 212 Apperson
Keiji Kanazawa, UC Berkeley

Reasoning about time and change is an important aspect of reasoning under uncertainty. It is essential to reasoning about plans and actions; it is also important in diagnosis, scene understanding, user modeling, and natural language understanding. This is especially true when such problems involve predictions of the behavior of a system based on observations of the system. In this section, we will study representations of change under uncertainty in a probabilistic framework. We introduce the study of stochastic processes, the mathematical theory of describing and predicting change. We describe Bayesian networks, logic, and other methods for representing stochastic processes and present some examples of their use in applications. We consider computational issues in using these representations. (Material in this section forms a basis for the section on planning and selective perception)

AM: 9:45 - 10:00 Break - 215 Apperson

AM: 10:00 - 12:00 Planning and Selective Perception - 212 Apperson
Thomas Dean, Brown

This lecture concerns the use of decision theoretic methods for automated planning and control under uncertainty. Markov decision processes are presented as a basic representation for planning under uncertainty. Standard notation and terminology is introduced from the literature on the control of stochastic processes. The lecture will describe a classical framework employing stochastic processes for decision making under uncertainty, and investigate how the sort of decomposable representations favored in artificial intelligence can be added to the framework to expedite inference. Existing approaches to decision-theoretic planning under uncertainty are surveyed and their relation to the classical framework is explained. It is claimed that Markov decision processes provide the semantic foundations for planning under uncertainty in much the same way as the propositional logic and its associated semantics provide the foundations for more expressive logics. This lecture focuses on the basic representational and computational issues of planning for Markov decision processes by making explicit the structure in such processes using graphical decision models. Applications from robotic control and airline scheduling are provided to ground the theoretical discussions.

12:00 - 1:30 Lunch, Gold Room, Arnold

PM: 1:30 - 1:45 Research Overview - 212 Apperson
Eric Horvitz, Microsoft Research

PM: 1:45 - 3:15 Abstraction - 212 Apperson
Mike Wellman, Univ. Michigan

In this session we will discuss the many ways that abstraction can be exploited in probabilistic reasoning. Generally speaking, the purpose of abstraction methods is to selectively ignore detail in order to save effort in reasoning or representation. There are several ways that we can do this in a probabilistic network model:

1. Abstracting probabilities (e.g., qualitative relationships, intervals)
2. Abstracting state spaces of random variables
3. Abstracting variable definitions (e.g., taxonomies of concepts)
4. Abstracting network structure (e.g., by ignoring weak dependencies)

Each method comes with costs and benefits, which may depend on context. In the tutorial session, we will review the abstraction methods that researchers have explored to date, and present some general issues bearing on any approach to abstraction in probabilistic reasoning.

PM 3:15 - 3:30 Break - 215 Apperson

PM: 3:30 - 5:30 Causation, Actions, and Qualitative Belief - 212 Apperson
Judea Pearl, UCLA

Bayesian belief networks are often defined as carriers of conditional independence information, while the causal interpretation of these networks is viewed as an optional bonus, or a curious side-effect. This tradition is changing rapidly. To better support practical decision making, current trends aim to base the theory of belief networks directly on the causal component. The result is a more natural understanding of what the networks stand for, what judgments are required in constructing the network and, most importantly, how actions and plans are to be handled within the framework of standard probability theory.

This tutorial will summarize the basic concepts in the new framework of "causal networks". Starting with functional description of physical mechanisms. we will derive the standard probabilistic properties of belief networks and show, additionally:

- * how the effects of propositionally-specified actions can be predicted from the network topology,
- * how qualitative causal judgments can be integrated with statistical data,
- * how persistence assumptions can be encoded in dynamic systems,
- * how actions interact with observations,
- * how counterfactuals sentences can be formulated and computed, and
- * how order-of-magnitude abstractions can yield a semi-qualitative decision theory.

PM 7:00 - 9:00 Dinner: Gold Room, Arnold
Perspective, Judea Pearl: A Journey into Neighboring Territories

Bexell lab open 8am - 5pm and 7pm - 10 pm
CS lab open 8am - midnight

Wednesday, 7/27

AM: 7:00 - 8:00 Breakfast: Gold Room, Arnold

AM: 8:30 - 10:15 Belief networks and learning - 212 Apperson
 Wray Buntine, NASA
 David Heckerman, Microsoft Research
 Stuart Russell, UC Berkeley

This talk will have two sections. There is too much material to cover any one thing in detail. This talk will instead give simple examples, and refer to suitable literature.

The first section will be a street-wise review of learning. Learning is an enormous area that, to the outsider, appears to have some four large but distinct communities, and six distinct theories, each addressing roughly the same tasks but with conflicting claims about their competing capabilities. The review will attempt to introduce the major issues and communities involved to help one navigate through the learning jungle. The review will also introduce some useful software.

The second section will be an introduction to learning with graphical models. We can model many of the classical learning problems from statistics, AI and neural networks with graphical models (basically, influence diagrams with undirected arcs and deterministic nodes), and many of the existing methods for learning can be summed up as combinations of a small set of standard exact and approximate graphical operators to simplify the problem: including Monte carlo sampling, exact methods, Laplace's approximation, EM algorithms, and differentiation operators. This section will introduce graphical methods for modeling standard learning problems, and then review the algorithms available. This section presents learning as an engineering problem rather than a research problem, and empowers one to rapidly prototype new algorithms for novel problems.

AM: 10:15 - 10:30 Break - 215 Apperson

AM: 10:30 - 12:00 Belief networks and learning, cont'd - 212 Apperson
 Wray Buntine, NASA

Noon: 12:00 - 1:30 Lunch, Gold Room, Arnold

PM: 1:30 - 2:30 Model Confidence - 212 Apperson
Kathryn Laskey, George Mason Univ.

PM: 2:30 - 3:30 Real-time and Embedded Systems - 212 Apperson
Bruce D'Ambrosio, Oregon State University
Eric Horvitz, Microsoft Research

Eric Horvitz will discuss Bayesian methods for decision making in time-critical contexts and describe several applications. The presentation will include issues surrounding the assessment and use of representations of time-dependent utility and a review of the use of decision-theoretic techniques for guiding computation and making decisions about tradeoffs in time-pressured contexts. The talk will include discussion of the role of flexible computational strategies for deriving action in situations of uncertain and varying time constraints. I will highlight key concepts in terms of several systems employing notions of time-dependent utility, including the Protos prototype for intensive care medicine and the Vista system for NASA Mission Control.

Bruce D'Ambrosio will very briefly sketch work on the OLMA system for real-time diagnosis, highlighting new experimental results confirming the utility of scenario-based anytime decision policy determination.

PM: 3:30 - 3:45 Break - 215 Apperson

PM: 4:00 - 5:00 Search and Scheduling - 212 Apperson
Othar Hansson, Heuristicrats Research

Because heuristic search is a completely general problem-solving method, the only limit on its applicability is the efficiency of search algorithms. That is a considerable limit, but as I will show, search is "simply" a problem of decision-making under uncertainty. One route to more efficient and predictable search algorithms is to design them to explicitly formulate and solve decision problems in the course of search.

We will study how decision theory and probabilistic inference are applied in two search applications: the Eight Puzzle and constraint-based scheduling of NASA telescopes. We will then discuss open research problems, related work, and the alarming correspondence between existing search algorithms and popular non-probabilistic uncertainty formalisms.

PM: 5:00 - 5:15 Wrapup - 212 Apperson, Bruce D'Ambrosio

PM: 7:00 - 8:00 Dinner: Gold room, Arnold

6 Appendix II - Participant List

The following is a complete list of all institute attendees, both presenters and non-presenters.

Name	Affiliation	email address
----	-----	-----
Greg Arnold	Wright-Patterson AFB	garnold@mbvlab.wpafb.af.mil
Marco Barbarisi	US Navy Coastal Sys Div.	marco@email.ncsc.navy.mil
Mindy Bokser	Calera Recognition Systems	mindy@calera.com
Michael Braverman	UC Berkeley CS	braver@CS.Berkeley.EDU
Jack Breese	Microsoft Research	breese@microsoft.com
Wray Buntine	NASA Ames	wray@ptolemy.arc.nasa.gov
Enrique Cortes-Rello	US West Advanced Tech Grp	ecortes@advtech.uswest.com
Stewart Crawford	Colorado State Univ. CS	sgcraw@cs.colostate.edu
Bruce D'Ambrosio	Oregon State University	dambrosi@research.cs.orst.edu
Thomas Dean	Brown	tld@cs.brown.edu
Rajarshi Das	Santa Fe Institute	raja@santafe.edu
Soren Dittmer	Univ. Aalborg	dittmer@iesd.auc.dk
Mike Frank	MIT CS	mpf@medg.lcs.mit.edu
Roman Fresnedo	Boeing Computer Services	roman@espresso.rt.cs.boeing.com
Robert Fung	Inst for Dec Sys Research	fung@netcom.com
Robert Goldman	Honeywell	goldman@src.honeywell.com
Othar Hansson	Heuristicrats Research	hansson@heuristicrat.com
Jordan Hayes	Heuristicrats Research	jordan@heuristicrat.com
David Heckerman	Microsoft Research	heckerma@microsoft.com
Max Henrion	Lumina Decision Systems	henrion@lumina.com
Allan Hollander	UC Santa Barbara Geography	adh@geog.ucsb.edu
Michael Horsch	Univ. British Columbia CS	horsch@cs.ubc.ca
Eric Horvitz	Microsoft Research	horvitz@microsoft.com
David Hovel	Microsoft Research	davidhov@microsoft.com
Charles Kahn	Univ. Wisconsin Radiology	ckahn@mcw.edu
Michael Jordan	MIT Psychology	jordan@psyche.mit.edu
Exaud Koka	U. Missouri-Rolla Civil E.	koka@umr.edu
Keiji Kanazawa	UC Berkeley	kanazawa@cs.berkeley.edu
Jak Kirman	Brown CS	jak@cs.brown.edu
Sven Koenig	CMU CS	planning
Alexander Kozlov	Stanford CS	alexvk@loki.stanford.edu
Kathryn Laskey	George Mason	klaskey@gmuvmx1.gmu.edu
Michael Laviolette	Univ. Missouri-Rolla Math	laviolet@umr.edu
Jim Leonard	Wright-Patterson AFB	jleonard@mbvlab.wpafb.af.mil
Tod Levitt	IET and Stanford	levitt@iet.com
Tony Maida	U. of Southwest Louisiana	maida@cacs.usl.edu
Laura Martignon	U. Ulm (Germany) Math	laura@neuro.informatik.uni-ulm.de
Kurt McCall	NASA Marshall Elec. Pwr	mccall@marvin.msfc.nasa.gov.

Deborah McGuinness	AT&T Bell Labs	d1m@research.att.com
Rulane Merz	Hughes Information Tech	rmerz@redwood.hac.com
Mark Minardi	Wright-Patterson AFB	mark_minardi@qm.mbvlab.wpafb.af.mil
Stan Musick	Wright-Patterson Labs	Stan.Musick@wpafb.af.mil
Madhura Nirkhe	Univ. Maryland CS	madhura@cs.umd.edu
Peter Norvig	Sun Microsystems Labs	Peter.Norvig@East.Sun.COM
Gary Ogasawara	Lockheed AI Center	garyo@cs.berkeley.edu
Jon Oliver	Open University, UK	J.J.Oliver@open.ac.uk
Judea Pearl	UCLA	pearl@cs.ucla.edu
Barney Pell	NASA Ames AI Branch	pell@ptolemy.arc.nasa.gov
Mark Peot	Knowledge Industries	peot@rpal.rockwell.com
William Poland	Stanford EES	poland@leland.stanford.edu
Gregory Provan	Inst for Dec Sys Research	provan@camis.stanford.edu
Koos Rommelse	Microsoft Research	koosr@microsoft.com
Scott Roy	Stanford CS	hsr@CS.Stanford.EDU
Stuart Russell	UC Berkeley	russell@cs.berkeley.edu
Larry Rystrom	Univ. Washington Physics	rystrom@apl.washington.edu
Sandip Sen	Univ. Tulsa	sandip@kolkata.mcs.utulsa.edu
Ross Shachter	Stanford	shachter@camis.stanford.edu
Mehdi Shirazi	Wright-Patterson AFB	mehdi_shirazi@qm.mbvlab.wpafb.af.mil
Suzanne Stanford	George Mason Univ. CS	sstanfor@mason1.gmu.edu
Magnus Stensmo	Salk Institute	magnus@salk.edu
Irena Vainshtein	Litton Data Systems	irina@lcsipdp.litton.com
Michael Wagner	Univ. Pittsburgh MIS	mmw@alpha.med.pitt.edu
Michael Wellman	U. Michigan	wellman@engin.umich.edu
Tony Zawilski	Mitre	zawilski@mitre.org
Yilin Zhao	Motorola IVHS	yilin@ivhs.mot.com

I had intended to write earlier but got sidetracked—in any case, I wanted to let you know that I really enjoyed the summer institute. I've been involved with a large number of such affairs over the years, and this one was the best I've seen, consistently good both in terms of organization and quality of the presentations.

Thanks, Mike Jordan MIT

STINFO